



Computing Sector

# Twelve Month Progress Report

## “Run II Data Preservation Project”

---

Version 1.0

3/4/2014

CD-DocDB #5186

PREPARED BY:

Robert D. Kennedy, Joseph B. Boyd, Kenneth Richard Herner, Bodhitha Jayatilaka, Stephen P. White

CONCURRENCES:

---

Rob Roser

Project Sponsor

---

Date

# Table of Contents

<b>1. PROJECT OVERVIEW .....</b>	<b>1</b>
<b>2. PROJECT STATUS.....</b>	<b>1</b>
2.1. PROJECT SCOPE.....	1
2.2. PROJECT SCHEDULE.....	1
2.3. PROJECT RESOURCES.....	3
2.4. PROJECT COST .....	3
<b>3. PROGRESS ON 12 MONTH MILESTONES .....</b>	<b>4</b>
3.1. DATA MIGRATION .....	4
3.2. DECIDE ON VIRTUAL MACHINE USE.....	4
3.3. ADDRESS ONE OPEN ISSUE.....	5
3.4. INTERNAL MILESTONE: PILOT VERSION OF FUTURE SYSTEM .....	5
<b>4. PROGRESS SUMMARY .....</b>	<b>6</b>
<b>APPENDIX 1. OPEN ISSUES LIST .....</b>	<b>7</b>

## 1. Project Overview

The purpose of this project is to implement data preservation for the Tevatron Run II experiments. The Run II data sets are very unique data sample of 1 TeV protons colliding with 1 TeV anti-protons, and represent a considerable investment. It is unlikely another data sample like this will ever be collected again. This sample is in many ways complimentary to the LHC data sample of multi-TeV protons on multi-TeV protons. A number of important “legacy” Tevatron physics measurements have been identified. While the majority of existing analysis tools will be available and supported for five years after end of data taking, the datasets will have value in generating and checking physics results throughout another decade.

Data preservation means different things to different people. The approach for this project will be to:

- curate the data from CDF and D0 Run II (both simulated and actual) by continuously migrating the data to modern storage media,
- maintain the full infrastructure ability to generate new Monte Carlo samples, simulate them, reconstruct them, and process them, and
- maintain the ability to perform physics analysis on both the simulated and actual data.

With this, a complete physics analysis can be carried out on Run II data well into the future.

## 2. Project Status

### 2.1. Project Scope

The Project Charter v1.0 (CD-DocDB #5072, 5/7/2013) was approved by the Project Sponsor and defines the project scope. There have been no scope changes since that charter version. Note that the project scope for each Run II experiment varies according to the needs of each individual experiment, so some project scope items refer only to one experiment or the other.


### 2.2. Project Schedule

The Project Charter v1.0 (CD-DocDB #5072, 5/7/2013) defines the project time frame. The project has chosen to include an additional set of milestones at 18 months (ending June 30,2014) in order to press the early delivery of the Run 2 Data Preservation system, allowing sufficient time for testing, adjustment, and hand-off to operations of the system before the project closes out. The timeline that the project continues to work towards is a little more aggressive than the charter:

- Jan – Jun 2013: Evaluate current state and plan the R2DP system
  - July 2013: Assess 6 Month milestones
- July – Dec 2013: Implement the R2DP pilot system, address some open issues
  - First significant changes/deployments permitted by experiments in Aug/Sep 2013
  - Jan 2014: Assess 12 Month milestones
- Jan – Jun 2014: Gap-fit the pilot system, address all open issues, finalize R2DP system plan
  - July 2014: Assess Internal 18 Month milestones
- July – Dec 2014: Deploy the R2DP system for both experiments and close-out the project
  - Oct – Nov 2014: Assess 24 Month milestones and begin project close-out

Because of the distributed decision-making, broad scope, and under-defined deliverables of the project, we have chosen not to maintain a detailed MS Project schedule. Rather, we are applying a rolling wave approach to plan the goals for each 3 to 6 month period, and then allow each experiment team to

implement those goals in the manner appropriate for that experiment for the work in-scope accepted by that experiment. The high-level project timeline is shown in Figure 1 below. Detailed R2DP tasks are tracked on the project's Sharepoint site (<https://sharepoint.fnal.gov/project/DataPreservation/default.aspx>) in the "Data Preservation tasks" list.



**Figure 1: Timeline for Identified Run II Data Preservation Work with progress (green) as of 2/21/2014**

The data migration goals were achieved for the 12-month mark. For more details on the data migration status, please see Section 3.1 Data Migration.

The goal of the CDF Offline Release is to produce a frozen offline release natively built on SL6, with an SL5 build as an intermediate deliverable. Initial progress on the SL6 build appeared promising, so the scientific validation step on SL5 was deferred. However, issues arose with the SL6 technical validation, and offline leadership transitioned, which delayed validating the release past the planned goal of Jan 1, 2014. We modeled the SL6 validation progress in 3 parts: Part 1 achieved the base build by Sep 2013, Part 2 validated applications and resolved Root I/O issues by mid-Feb 2014, and finally Part 3 is addressing issues with the “maxopt” optimized build used in large-scale processing. Due to the extended SL6 validation timeline, the SL5 scientific validation was performed afterall. That SL5 build was validated and used in the CDF R2DP pilot. The SL6 build is expected to be validated by the end of February 2014.

The goal of the DZero Offline Release is to produce a frozen offline release built on SL6 using some SL5-compatibility libraries. Technical validation and scientific validation of the release were successful and completed before the end of January 2014. This release was used for the DZero R2DP pilot.

## 2.3. Project Resources

The Project Charter v1.0 (CD-DocDB #5072, 5/7/2013) defines the project resources and organization at a high level. There have been no major changes to the resources and organization since that charter version. The extended project team continues to evolve as the Run II experiment support is transitioning towards more efficient smaller-scale operations, which has led to more opportunities to align data preservation “operations” work with this project’s work.

## 2.4. Project Cost

The Project Charter v1.0 (CD-DocDB #5072, 5/7/2013) defines the project budget. There have been no changes to the budget since that charter version. The budget for FY13 is 3.97 FTEs in SWF and zero for M&S. In FY13, 2.98 FTEs of effort were charged to the project, which is 75% of the FY13 SWF budget. As such, the project is moderately underspent. We plan engage more resources on project work in the 12-18 month stage of the project as “operations” work on Run II experiments is becoming more aligned with this project’s work. The cumulative project effort to date for FY13 is shown in Figure 2 below.




Figure 2: Cumulative Effort Reported in FY13 for the Run II Data Preservation Project

### 3. Progress on 12 Month Milestones

#### 3.1. Data Migration

Goal: Complete 50% of data migration to modern T10K tape technology

Description: The project oversees this request to the Mass Storage Services (MSS) group, who will actually perform the actual work. The project and the MSS group work with the experiments to assess priorities in data migration, to insure that on-going data analysis takes priority over this data migration. In order to achieve the 50% data migration goal per 6 months of the project, each experiment data migration effort should sustain on average about 1% of the data migration per week.

Results: as of 1/6/2014

- CDF: 87.5% of data migrated
  - The CDF data migration continues to sustain an average data migration rate greater than 1% per week in this time period. If this rate continues, the CDF data migration may be completed in mid-2014. We are suggesting to the experiment to switch to T10K tape technology for data newly written to tape to allow the migration to complete smoothly.
- DZero: 50.2% of data migrated
  - The DZero data migration sustained a rate well below 1% for the first 6 months of the project. The MSS group investigated and found low-level disk contention in the mover nodes for the migration. After replacing disks on these nodes with solid-state disks in early August 2013, the DZero data migration sustained a period of greater than 1% migration rate through December 31, 2013.
  - Since January 1, 2014, tape system issues have slowed progress. This has caused the DZero data migration to fall behind pace for the 75% milestone as shown in Figure 1.

Assessment: **Goal achieved.**

Next Steps: Project should continue to track data migration rates to ensure continued success.

#### 3.2. Decide on Virtual Machine Use

Goal: Decide on whether to go to virtualized machines.

Description: The project has decided to move Run II services/systems to virtualized machines in a graded approach. The default choice will be to move services running on physical machines to run on virtual machines. However, some services running on physical machines that are nearing end-of-life and are not critical to experiment operations will be allowed to naturally migrate to virtual machines when the physical machine hosts are retired. We created tasks in the project timeline to track the migration of services to virtual machines for both experiments. Note that this topic is also treated in Open Issue 9a.

Results: The R2DP system is working with the Run II experiments and SCD to move to a virtual machine-based infrastructure. This work is well underway, with some services intentionally moving later to virtual machines with the refresh of their host physical hardware.

Assessment: **Goal achieved.**

Next Steps: We will track progress on these migrations via weekly project status reports.

### 3.3. Address One Open Issue

Goal: Address one of the Open Issues

Description: The project has identified Open Issues and a plan to resolve them. Open Issues are high-level unresolved issues worthy of reporting to the Project Sponsor. Some Open Issues are in fact risk mitigation believed to be necessary to ensure that the Run II Data Preservation system is sustainable by the Fermilab Scientific Computing Division with minimal effort from the Run II experiments.

Results: Appendix 1 contains the updated Open Issues list. Of the identified 11 high-level Open Issues:

- 3 are fully CLOSED.
- 4 are partly closed, with some planned work remaining.
- 4 are OPEN and have a plan and timeline defined.

Assessment: **Goal achieved.**

Next Steps: The project has begun executing the Open Issues plans. For Open Issues that are determined to not be actionable, the project will develop risk plans and support agreements that all parties (CDF, DZero, SCD) can accept.

### 3.4. Internal Milestone: Pilot Version of Future System

Goal: Develop and test an early “pilot” version of the future R2DP system for both experiments.

Description: In order to test some of the design of the Run II Data Preservation system as soon as possible and gauge progress, the project defined a pilot version with features that would be available in the 12 month timeframe. The available features varied by experiment, with one experiment taking the lead with some features, while the other experiment took the lead on others, and then each will be able to consult the other’s experience in the later stages.

Results:

- The CDF pilot system was successfully tested. Some of the differences between the pilot and final CDF R2DP systems are:
  - Used SL5-based CDF release instead of the future SL6-based CDF release.
  - Did not have IF-compatible job submission in place yet
  - Did not yet use SAMWeb yet.
  - Used dCache underneath SAM Station – same as has been in use at CDF.
  - Did not use CVMFS yet.
- The DZero pilot system was successfully tested. Some of the differences between the pilot and final DZero R2DP systems are:
  - Used SL6-built DZero release per plan.
  - Job submission script is not ready yet, but job submission scheme tested.
  - Used SAMWeb on CAB and on a FermiCloud node.
    - Also ran via FermiGrid without SAM (manual file delivery).
  - Emulated dCache underneath SAM Station, tested on a FermiCloud node.
    - dCache underneath SAM has since been tested and works (February 2014).
  - Used CVMFS.

**Assessment: Goal achieved.**

**Next Steps:** The project is defining a quarterly release plan to pace the introduction of the features missing in the pilot system to achieve the final R2DP configurations in mid-2014. That plan is on the project's Sharepoint site (<https://sharepoint.fnal.gov/project/DataPreservation/default.aspx>) in the "Data Preservation tasks" list.

## **4. Progress Summary**

Overall, the R2DP project continues to achieve its goals. Tape data media migration continues at an acceptable pace. Though CDF's frozen offline release has taken longer to validate than expected, the timeline of the project has not been affected yet since a SL5 validated release was available. The CDF Offline database upgrade was a success, and Frontier has moved to virtual machines with a long-term support plan. DZero RD2P leveraged work by DASPOS resource to adapt to CVMFS use. Progress was made setting up a DZero dCache to allow testing of dCache underneath SAM at DZero. Planning has begun for a more detailed feature delivery schedule as both experiment's efforts have achieved success in the initial pilot. We expect this success to continue as the project works towards a tested initial release of R2DP system for both experiments by July 1, 2014.



## Appendix 1. Open Issues List

We have grouped the Open Issues into two main categories, issues from the users' point of view and issues from the service providers' point, to reflect our goal of delivering a usable and sustainable system. Documentation references are to the R2DP Project site and require Fermilab SERVICES login.

To date, 3 of 11 Open Issues are closed and 4 of 11 Open Issues are partly closed.

### How will data analysis be performed in the Run II Data Preservation system? – users' point of view

1. How are we going to perform bookkeeping?
  - a. Plan: Both Run II experiments will move from using SAM to using SAMWeb for a sustainable data bookkeeping solution.
  - b. CDF Documentation: [SAMWeb for CDF \(Illingworth\).pdf](#)
  - c. DZero Documentation: [SAMWeb for DZero.pdf](#)
  - d. Status: OPEN
  - e. Timeline: Project will close this by 18-month milestone.
2. How will a user be authorized to access data?
  - a. Answer: The R2DP system will continue with the existing experiment management plan for conceptual authorization even after the Run II experiments end.
    - i. This depends on some informal Run II collaboration existence that can define who is a member of the "experiment" and thus qualified to access the data.
    - ii. The technical means of authorization will remain the same as well.
  - b. Status: CLOSED (6-month)
3. How will jobs be run? How will the user start his/her job?
  - a. Plan: Submit jobs using the Intensity Frontier (IF) job submission system.
    - i. DZero Plan: We plan to integrate existing DZero job submission scripts with the IF system to hide this change from the end-users.
    - ii. CDF Plan: We plan to use the same CAFsubmit. Nothing changes for end-users.
  - b. Plan: Jobs will run on a grid or cloud node. Details of this are under discussion.
  - c. Status: OPEN
  - d. Timeline: Project will close this by 18-month milestone provided the IF job submission system work is completed in time (outside project's control), else a work-around will be put in place in preparation for the IF job submission release by 24-month milestone.
4. Is our MC framework flexible enough to interface with new generators?
  - a. Answer: Yes, because the input format to the MC framework can always be flat ASCII text file.
  - b. Status: CLOSED (6-month)
5. Can we use new data analysis programs as they become available?
  - a. Answer: This is up to the user to test and implement, not the project. Past experience with ROOT encourages us to believe that users will be able to use new versions of ROOT on existing data either directly or after applying a ROOT-supplied conversion tool.  
Note:
    - i. New versions of ROOT will not be centrally managed or deployed by project.

- ii. Experiment code will not be built and validated against new versions of ROOT by the project.

**b. Status: CLOSED (6-month)**

6. How will users be supported for data analysis questions?

- a. Answer: While not strictly in project scope, perceived success of system depends on this. The elements the project may pursue to encourage this outcome include the following:
  - i. Identify and address gaps in support process due to fewer active users and analyses. This is an Experiment task. The Project is a stakeholder, but not responsible for execution.
  - ii. Update user documentation known to be out-of-date. This is an Experiment task. The Project is a stakeholder, but not responsible for execution.
  - iii. User analysis support system – there should be reduced expectations for support response time, and greater reliance on documentation. This is an Experiment task. The Project is a stakeholder, but not responsible for execution.
- b. Plan: The Project will help identify areas where data analysis support may be weak to the experiment as we enable example analyses to run in the Run II Data Preservation system, to ensure the perceived quality of the project deliverable is not reduced by lack of analysis support.

**c. Status: OPEN**

- d. Timeline: Project will close this by 18-month milestone.

**How will the infrastructure be maintained for a Post-2015 system? – service provider's point of view**

7. What needs to be upgraded to current versions to re-establish version currency?

- a. Data migration to modern T10k media
  - i. Plan: track data migration. Work with the Mass Storage group and the experiments to address bottlenecks in the migration.
  - ii. Timeframe: This migration will continue until at least July 1, 2014.
  - iii. **Status: OPEN**
  - iv. Migration Progress Tracking:

<https://sharepoint.fnal.gov/project/DataPreservation/SitePages/TapeMigration.aspx>

- b. CDF: Offline Oracle RDBMS to 11
  - i. Plan: upgrade the Offline development RDBMS first and test operation. Then, after the summer 2013 conferences, upgrade the Offline production RDBMS.
  - ii. Update: Upgrade was successfully performed 9/10/2013.
  - iii. **Status: CLOSED (12-month)**
- c. CDF: Frontier/Squid
  - i. Plan: upgrade to the most modern version of Frontier and its underlying toolset that are usable by CDF.
  - ii. Update: All servers updated in July 2013. Support is transitioning from CDF experimenter to SCD (with experimenter consulting) per agreement.
  - iii. CDF Documentation: [CDF Frontier.docx](#)
  - iv. **Status: CLOSED (12-month)**
- d. **Overall Status: OPEN**
- e. Timeline: Project will close by the 24-month milestone.

8. What should be changed/upgraded to reduce the risk of obsolescence in existing infrastructure?
  - a. Convert to using new SAM server
    - i. Plan: Both Run II experiments will move from using SAM to using SAMWeb for a sustainable data bookkeeping solution. DZero has basically already completed this process, though SAMWeb is not yet the default.
    - ii. CDF Documentation: [SAMWeb for CDF \(Illingworth\).pdf](#)
    - iii. DZero Documentation: [SAMWeb for DZero.pdf](#)
    - iv. Status: Open
  - b. DZero: Calib DB Servers
    - i. Plan: Deploy a virtual node with the new SLF6-based releases of all Calibration DB servers. This will replace the Calibration DB servers currently on d0dbsrv20. We plan to execute this move after the Summer 2013 conference season.
    - ii. Note: This only partially overlaps Issue (9a) since the service software is also being updated and ported to work on a newer operating system.
    - iii. Update 2/2014: One Calib DB Server upgraded and deployed. Remainder are ready to be deployed at next regularly scheduled downtime.
    - iv. Status: Open.
  - c. Migrate logbooks to supported tools
    - i. CDF Status: already using the supported tool.
    - ii. DZero Status: Migration has been converted to ECL.
      1. See <http://dbweb0.fnal.gov/ECL/dzero>
    - iii. Status: Closed (6-month).
  - d. Migrate some auxiliary databases, database applications to supported tools
    - i. CDF Plan: People DB does not have a plan yet. Internal Notes going to Spires.
    - ii. DZero Plan: Whod0, Speakers Bureau will be dumped into a simple archive format. Internal Notes migrated to Spires.
    - iii. Status: Open
  - e. DZero: Transition from CORBA to HTML
    - i. Plan: Transition plan documented, but experiment chose not to proceed
    - ii. Documentation: DZero Documentation: [D0 Calibration DbServers](#), [D0 CORBA Emails](#), [CORBAtoHTTP Effort Summary](#)
    - iii. Agreement: We have met with the DB group and the D0 algorithms group and agreed that the html transition will occur if and when CORBA stops working. CORBA infrastructure will not be fixed. D0 understands that the CORBA-to-HTTP transition may take longer to complete if the work is done later rather than now owing to experts potentially moving on to other projects. D0 also understands that MC production will be unavailable for the entire CORBA-to-HTML transition period. D0 accepts the risks of not performing the transition now in order to save valuable manpower for other projects during the 2013-14 timeframe.
    - iv. Status: Closed (6-month).
  - f. How long are certain production servers needed, before being retired?
    - i. Plan: We are collecting a comprehensive list of candidate production servers. Many have already been migrated as driven by the refresh cycle or risk issues.
    - ii. Documentation: [DPPlan.docx](#)
    - iii. Status: Open
  - g. Overall Status: OPEN
  - h. Timeline: Project will close this by 18-month milestone.

9. How/where will we adapt to enterprise architecture guidance to reduce operations costs?
  - a. Move services on “physical nodes” to virtual nodes (tied to Issue 8f)
    - i. Plan: These transitions are on-going, and will continue until all services for the Run II Data Preservation system that can be migrated to virtual nodes are.
    - ii. **Status: Open.** Much has been done (Frontier, CalibDB Servers), but working on comprehensive plan for all services.
    - iii. Documentation: [DPPlan.docx](#)
  - b. Move services on experiment-specific systems to shared systems where possible
    - i. Plan: List which services can be moved/shared and transition accordingly.
    - ii. **Status: Open.** Plans are being developed for larger such systems such as Mass Storage which may or may not be moved to shared systems/nodes.
  - c. **Overall Status: OPEN**
  - d. Timeline: Project will close this by 18-month milestone.
10. How is this system going to be technically sustainable from 2015 through 2020? Roadmap?
  - a. SLF6 versus SLF7
    - i. Plan: Focus on an all-SLF6 system that operates in a virtual environment. The host CPUs may run a different OS, but the analysis environment will be SLF6.
    - ii. Documentation: [D0 s/w beyond 2020.pdf](#), [CDF task force report](#), [DPPlan.docx](#)
    - iii. **Status: Closed (6-month).**
  - b. Oracle RDBMS
    - i. Plan: Investigate the potential roadmap for Oracle. Develop a risk plan that takes into account the cost of upgrading versus not-upgrading especially given recent positive experience operating on unsupported Oracle RDBMS versions at CDF.
    - ii. Documentation: [DPPlan.docx](#)
    - iii. **Status: Closed (6-month).**
  - c. DZero: SAM Cache and dCache (for local data caching)
    - i. Plan: Transition to using dCache for this role, a product with long-term support.
    - ii. Documentation: [Dzero Job Submission Talk](#)
    - iii. **Status: Open.**
  - d. **Overall Status: OPEN**
  - e. Timeline: Project will close this by 18-month milestone.
11. How is this system going to be financially sustainable from 2015 through 2020? Support plan?
  - a. Determine how to prove the R2DP system is operational after changes.
    - i. Plan: Identify and document official validation test procedures, standard analyses
    - ii. Plan: Determine if automating some or all validation testing is feasible.
    - iii. **Status: Open.**
  - b. CDF: Frontier/Squid Support
    - i. Plan: Develop a plan to migrate support responsibility from CDF to SCD.
    - ii. Status: Squid support has not yet been formally worked out with SCD.
    - iii. **Status: Closed (12-month).**
  - c. How will costs for a possible future media migration be addressed?
    - i. Plan: engage SCD management on this question. This is heavily dependent on future CMS choices of media, equipment, and robotics.
    - ii. **Status: Open.**
  - d. **Overall Status: OPEN**
  - e. Timeline: Project will close this by 18-month milestone.